

In the Claims

Please amend the claims as follows.

1. (Amended) A method of estimating the results of a database query, the method comprising:
 - collecting workload information related to queries that have been executed on the database;
 - tracing query patterns of the queries in the workload to identify the usage of tuples in the database during execution of the queries wherein the usage of a given tuple relates to a relative frequency with which the given tuple was accessed by the queries in the workload;
 - determining a sample weights based on tuple usage for each tuple; and,
 - performing a weighted sampling of the database based upon the sample weights; and
 - executing the database query on the weighted sample to estimate results of the database query.
2. (Amended) The method of claim 1 wherein the weighting weighted sampling is performed by assigning a weight to each tuple based on a probability of usage of the tuples required in executing the queries in the workload.
3. (Original) The method of claim 2 and further comprising computing an aggregate over values in each sample tuple.
4. (Original) The method of claim 3 wherein the aggregate is computed by multiplying each value by the inverse of the probability with which corresponding tuples were sampled.
5. (Original) The method of claim 1 wherein the weights are a function of the frequency of access of a tuple and the number of queries in the workload that access the tuple.
6. (Original) The method of claim 1 wherein the tuple usage is stored on a page level.

7. (Amended) A machine readable medium having instructions for causing a machine to perform a method of estimating the results of a database query, the method comprising:

collecting workload information related to queries that have been executed on the database;

tracing query patterns of the queries in the workload to identify the a usage of tuples in the database during execution of the queries wherein the usage of a given tuple relates to a relative frequency with which the given tuple was accessed by queries in the workload;

determining a sample weights based on tuple usage for each tuple; and,

performing a weighted sampling of the database based upon the sample weights; and

executing the database query on the weighted sample to estimate results of the database query.

8. (Original) The machine readable medium of claim 7 wherein the weights are a function of the frequency of access of a tuple and the number of queries in the workload that access the tuple.

9. (Original) The method of claim 7 wherein the tuple usage is stored on a page level.

10. (Amended) A system that estimates the results of a database query, the method comprising:

a module that collects workload information related to queries that have been executed on the database;

a module that traces query patterns of the queries in the workload to identify the usage of tuples in the database during execution of the queries wherein the usage of a given tuple relates to a relative frequency with which the given tuple was accessed by queries in the workload;

a module that determines a sample weights based on tuple usage for each tuple; and,

a module that performs a weighted sampling of the database based upon the sample weights ; and

a module that executes the database query on the weighted sample to estimate results of the database query.

11-34. (Canceled)

35. (Amended) A method of estimating the results of a database and a given workload wherein the queries in the workload may have selection conditions, the method comprising:

collecting workload information related to queries that have been executed on the database;

tracing query patterns of the queries in the workload to identify the a usage of tuples in the database during execution of the queries wherein the usage of a given tuple relates to a relative frequency with which the given tuple was accessed by queries in the workload;

determining a sample weights based on tuple usage for each tuple;

performing a weighted sampling of the database based upon the sample weights;

executing the database query on the weighted sample to estimate results of the database query, and,

generating a weighted outlier index.

36. (Original) The method of claim 35 and further comprising calculating an aggregate based on the samples of the index.